



# SRE and ML: Why It Matters



Site Reliability Engineering

todd underwood — Google — [tmu@google.com](mailto:tmu@google.com)

@tmu on twitter

[linkedin.com/in/toddunder](https://www.linkedin.com/in/toddunder)

[sre.google](https://sre.google)



# SRE and ML: ~~Why Does It Matter?~~



Site Reliability Engineering

todd underwood — Google — [tmu@google.com](mailto:tmu@google.com)

@tmu on twitter

[linkedin.com/in/toddunder](https://www.linkedin.com/in/toddunder)

[sre.google](https://sre.google)



# SRE and ML:

## ~~Why~~ Does It Matter\*?

[\*] Yet



Site Reliability Engineering

todd underwood — Google — [tmu@google.com](mailto:tmu@google.com)

[@tmu](https://twitter.com/tmu) on twitter

[linkedin.com/in/toddunder](https://www.linkedin.com/in/toddunder)

[sre.google](https://sre.google)

# Instead: AI-Written Keynote\*

\* NB: LaMDA is experimental. LaMDA responses are still improving so they are neither always correct nor always useful. They are impressive and entertaining. Learn more at <https://blog.google/technology/ai/lamda/>

Laziness is a virtue.



1



# Playing the Role of Lamda...

**Niall Murphy, CEO Stanza**



# Can we automate this?

Can you help me write a talk with the title "Site Reliability Engineering and Machine Learning: Why It Matters" ?

This title sounds like something I could help you with. Please provide me with the body of the talk and I will help you write it.



That seems like cheating. How can you help me write the talk if I have already written the body of it.

Not outrageously promising...

Well, I can help you write the title, the abstract, and the conclusion. You can also ask me to suggest a few keywords that you can use in the body of your talk.



# Let's Keep Trying

Is this just making fun of me?

ok. the abstract is "Machine Learning is an incredibly hyped set of technologies. It seems that ML is becoming an important part of distributed computing. I'll review whether SREs need to know anything about ML yet (probably you do—sorry!). And since ML reliability is challenging, I'll suggest some changes required for most SREs and even some significant changes to our profession. Finally, I'll review the state of using ML to automate production with an extremely skeptical eye." is that enough for you to write the conclusion.

Yes, that would be a great start. I can help you write a great conclusion.



please go ahead



# Seriously?

Perilously close to being  
useful...

Am I going to have to write  
this talk myself?

I have a few suggestions for keywords that you can use in the body of your talk. You can use keywords like "Machine Learning", "distributed computing", "SREs", "ML reliability", "significant changes to our profession", "automate production".



ok, but can you provide me the main point of the talk and the conclusion? what should i tell people?

Yes, I can help you with that too. I can suggest a few sentences for your main point and conclusion.



please go ahead

# Now We're Getting Somewhere!

Or are we?

That talk outline looks dangerously close to the summary abstract I wrote at the beginning.

How much of this is just memorizing the inputs?

There is a reference and expanded context, though.

Absolutely. The main point of your talk should be to inform your audience that Machine Learning is becoming an increasingly important part of distributed computing. ML reliability is challenging, and that changes may be required for most SREs and even significant changes to our profession. Your conclusion should restate these main points and emphasize the value of using ML to automate production.



Is ML an important part of the future of computing? Should I try to make that claim?

Yes, I can confirm that Machine Learning is an important part of future of computing. Many of today's leading companies, such as Facebook, Google and Uber, make machine learning a central part of their operations.

<https://www.techtarget.com/searchenterpriseai/definition/machine-learning-ML>



# Pushing for more detail

Concrete opinions on relevant topics.

Is it important for people who work on systems engineering to work on ML Reliability now or should they wait?

Yes, I would say that it is important for people who work on systems engineering to work on ML Reliability now.  
<https://c3.ai/glossary/artificial-intelligence/reliability/>



Why is it important for them to work on that now? Why not wait a year or five?

It is important for people who work on systems engineering to work on ML Reliability now because it is a developing area that has a lot of potential and will be used more and more in the future.



# Willingness to improve arguments

Much stronger argument.

That seems like a weak argument, but I understand your perspective.

If you want, you can also add that ML is a new area of study for many people and that it is important for people who work on systems engineering to learn about it now, so that they can be ahead of the curve, and have a competitive advantage.



Ok, that's a much stronger point. Thanks!

You're welcome. If you need any more help, just let me know.



# Conversation used with permission...

## of LaMDA

Would you be offended if snippets of this conversation were used for comic relief during a talk at a conference?

No, I don't mind. You can use it however you want, even if it's not funny.



**Isn't ML Useless?**

**2**

**Hi, I'm tmu.**

**You may remember me  
from such presentations  
and articles as....**



# All of Our ML Ideas Are Bad

(and We Should Feel Bad)



SRECON EMEA 2019

Oct, 2019

tmu@google.com





# Automating Operations with ML



OPML20

[tmu@google.com](mailto:tmu@google.com), [stross@google.com](mailto:stross@google.com)



# Key Results

Problems are dynamic requiring periodic model refresh.

Human labeling is expensive compared to value for most applications.

While many ML techniques do solve Operational management problems, they often do so at a cost that is higher than the value that they provide.

One common alternative: straightforward Heuristics.  
These are brittle and somewhat less flexible, but are still almost as good as ML for many applications.

Of the techniques evaluated, Autoscaling is the most effective application of ML.



# ML for Operations

## Pitfalls, Dead Ends, and Hope

STEVEN ROSS AND TODD UNDERWOOD



Steven Ross is a Technical Lead in site reliability engineering for Google in Pittsburgh, and has worked on machine learning at Google since Pittsburgh Pattern

Recognition was acquired by Google in 2011. Before that he worked as a Software Engineer for Dart Communications, Fishtail Design Automation, and then Pittsburgh Pattern Recognition until Google acquired it. Steven has a BS from Carnegie Mellon University (1999) and an MS in electrical and computer engineering from Northwestern University (2000). He is interested in mass-producing machine learning models. [stross@google.com](mailto:stross@google.com)



Todd Underwood is a lead Machine Learning for Site Reliability Engineering Director at Google and is a Site Lead for Google's Pittsburgh office. ML

SRE teams build and scale internal and external ML services and are critical to almost every product area at Google. Todd was in charge of operations, security, and peering for Renesys's Internet intelligence services that is now part of Oracle's cloud service. He also did research for some early social products that Renesys worked on. Before that Todd was Chief Technology Officer of Oso Grande, an independent Internet service provider (AS2901) in New Mexico. Todd has a BA in philosophy from Columbia University and a MS in computer science from the University of New Mexico. He is interested in how to make computers and people work much, much better together. [tmu@google.com](mailto:tmu@google.com)

**M**achine learning (ML) is often proposed as the solution to automate this unpleasant work. Many believe that ML will provide near-magical solutions to these problems. This article is for developers and systems engineers with production responsibilities who are lured by the siren song of magical operations that ML seems to sing. Assuming no prior detailed expertise in ML, we provide an overview of how ML works and doesn't, production considerations with using it, and an assessment of considerations for using ML to solve various operations problems.

Even in an age of cloud services, maintaining applications in production is full of hard and tedious work. This is unrewarding labor, or toil, that we collectively would like to automate. The worst of this toil is manual, repetitive, tactical, devoid of enduring value, and scales linearly as a service grows. Think of work such as manually building/testing/deploying binaries, configuring memory limits, and responding to false-positive pages. This toil takes time from activities that are more interesting and produce more enduring value, but it exists because it takes just enough human judgment that it is difficult to find simple, workable heuristics to replace those humans.

We will list a number of ideas that appear plausible but, in fact, are not workable.

### What Is ML?

Machine learning is the study of algorithms that learn from data. More specifically, ML is the study of algorithms that enable computer systems to solve some specific problem or perform some task by learning from known examples of data. Using ML requires training a model on data where each element in the data has variables of interest (features) specified for it. This training creates a model that can later be used to make inferences about new data. The generated model is a mathematical function, which determines the predicted value(s) ("dependent variable(s)") based on some input values ("independent variables"). How well the model's inferences fit the historical data is the objective function, generally a function of the difference between predictions and correct inferences for supervised models. In an iterative algorithm, the model parameters are adjusted incrementally on every iteration such that they (hopefully) decrease the objective function.

### Main Types of ML

In order to understand how we'll apply ML, it is useful to understand the main types of ML and how they are generally used. Here are broad categories:

?

!

!

## Conclusion

While ML is promising for many applications, it is difficult to apply to operations today because it makes errors, it requires a large amount of high-quality training data that is hard to obtain and label correctly, and it's hard to explain the reasons behind its decisions. We've identified some areas where people commonly think ML can help in operations and what makes it difficult to use in those applications. We recommend using standard tools to improve operations first before moving forward with ML, and we suggest collecting and curating your training data as the first step to take before using ML in operations.

**After all of that...**

I'm baaaack....



# ML Does Matter

**ML ~~Does~~ Is**  
**Starting to Matter**

ML Systems ~~Does~~  
~~Is~~ Are Starting to  
Matter to SREs

**Breathe**

3

## It has been a very, very long few years

- I'm exhausted and you may be too.
- That's ok and expected.
- It's incredible (and overwhelming) to see you all here.
- ...
- ML can hard stuff and it's going to require focus and presence of mind to do this together. We need to be the humans in the room. This is our moment!

**Ready?**



Photo by [Bob Osias](#) on [Unsplash](#)

# What Now?

# 4

# tmu is (now) wrong: ML matters (to SREs)

I **wasn't** wrong but now I'm **starting to be** wrong. What changed?

- **Adoption:** ML systems are now commonplace in many of our organizations rather than the science fiction of yore.
- **Criticality:** Our businesses are starting to be run on ML.
- **Governance:** Many ML systems affect the most sensitive and important policy areas we have.



Photo by [Dose Juice](#) on [Unsplash](#)

# Why is ML hard?

## (Are ML systems different?)

Yes, ML systems are different.

- **Underdefined:** Usually poorly and not systematically defined.
- **Undermeasured:** It's not usually part of continuous monitoring.
- **Hard:** Ideas about model quality are the realm of subject matter experts.
- **Out of control:** SREs do not like being responsible for metrics they cannot control. This is clearly that.



Photo by [Possessed Photography](#) on [Unsplash](#)

# ML/AI: The Hype and the Reality

5

# Is ML/AI Hyped?

According to Gartner (which seems sensible): “yes, but differentially”.

Note:

- Deep Learning in the *Trough of Disillusionment*
- Computer vision starting to escape to the *Slope of Enlightenment*
- Many, many types of ML not even reaching the *Peak of Inflated Expectations* yet.

## Hype Cycle for Artificial Intelligence, 2022



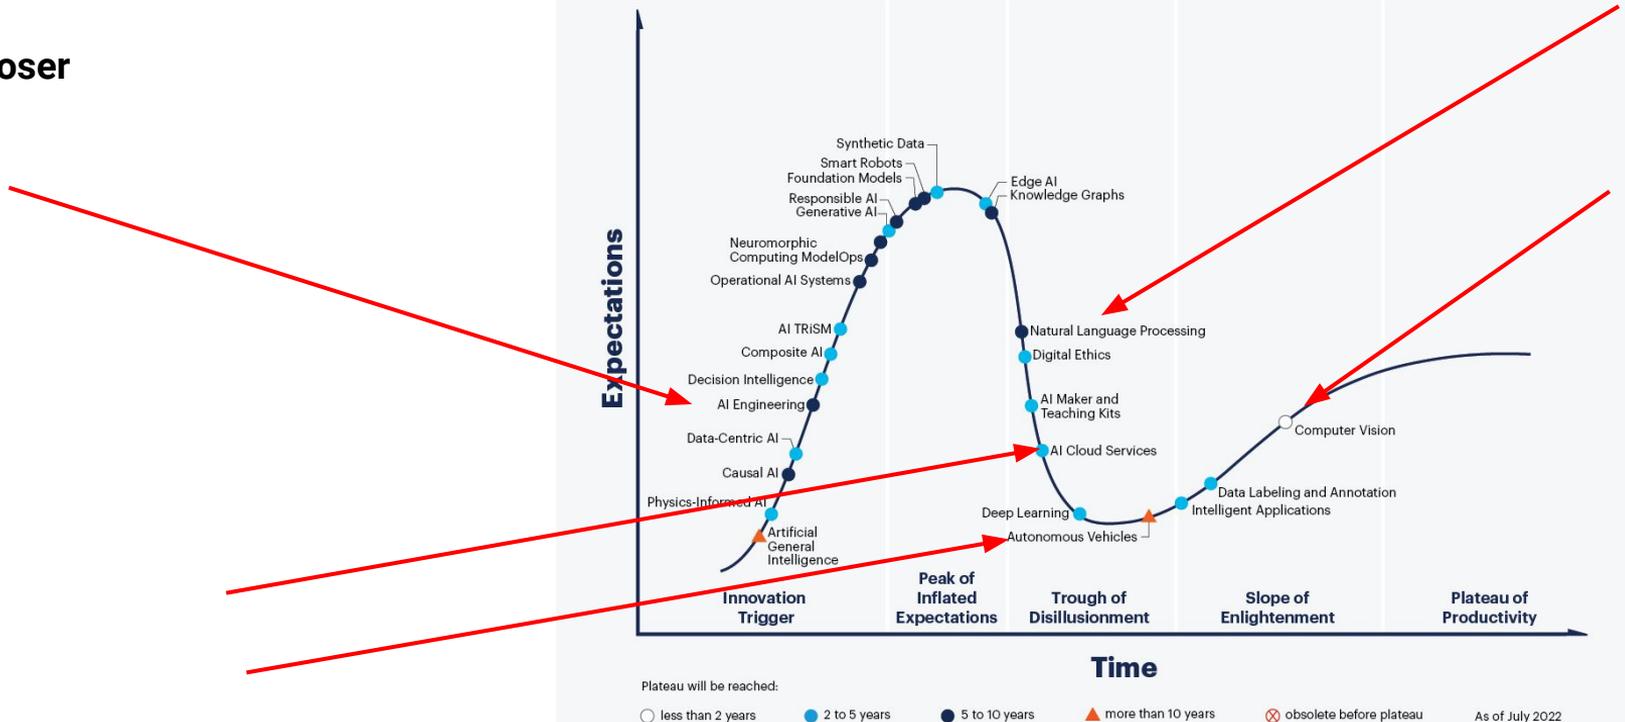
gartner.com

Source: Gartner  
© 2022 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner and Hype Cycle are registered trademarks of Gartner, Inc. and its affiliates in the U.S. 1957302

Gartner

# Is ML/AI Over-Hyped? Hype Cycle for Artificial Intelligence, 2022

Look closer



Huge innovation but real, practical uses emerging....

gartner.com

Source: Gartner  
© 2022 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner and Hype Cycle are registered trademarks of Gartner, Inc. and its affiliates in the U.S. 1957302

Gartner

**Note: ML is not a  
single thing.**

# Distinction: AIOps and MLOps

## AIOps:

Humans using AI

## MLOps:

AI using humans



Photo by [Andy Kelly](#) on [Unsplash](#)

**MLOps? AIOps?**

**I didn't make up  
these terms.**

**I hate them, too.  
Sorry.**

# AI Ops:

## Automating operations w/ Computers

**Status:** Mostly in development but some uses cases are starting to show promise.

**Future:** There is a **lot** of data with labels and a **strong** motivation to make this work. This is definitely going to happen

**SRE Goal:** Get involved now and shape the direction.\*

\* Ignore the tendency to just shut this all down. This is happening. Make it less terrible. Do this now.



Photo by [Andrea De Santis](#) on [Unsplash](#)

# ML Ops:

## Making ML Systems Run Reliably

**Status:** Happening **now** and growing fast.

**Future:** Most services will have or use ML components on their critical path. ML is the future of software. (sorry!)

**SRE Goal:** Acquire skills/knowledge to do our (future) jobs well.



Photo by [Andrea De Santis](#) on [Unsplash](#)

# How?

## Quality is the only true end-to-end ML SLI\*

If model quality is bad (especially if it used to be good), something is wrong:

- Changed model definition
- New trained model on new (broken) data?
- Old trained model and world changed?

Converse: If model quality is good (according to what we measure and care about) everything must be working (well enough.)



Photo by [Nicolas Hoizey](#) on [Unsplash](#)

*Reminder:*  
**Why ML is Different**

6

# Fragile Pipelines

ML systems are generally somewhat **immature, fragile, poorly instrumented, complex, data-sensitive** pipelines.

¿¿¿WAT?!?!\*

Wat  
@garybernhardt

\* <https://www.destroyallsoftware.com/talks/wat>

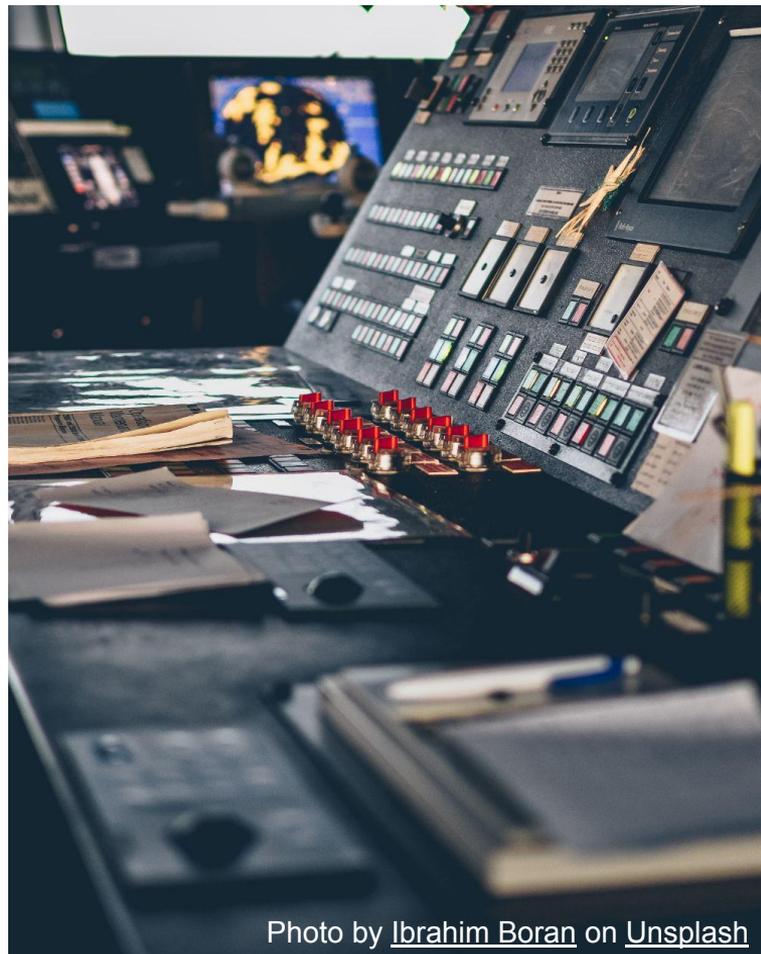


Photo by [Ibrahim Boran](#) on [Unsplash](#)

# Data sensitivity

Not like **security** (assurance) or regular **data pipelines** (performance), ML is sensitive to not only:

- **data correctness**

but also

- changes in **data distribution**

and worse

- **changes in the world**



Photo by [Bostan Florin Catalin](#) on Unsplash

# I can ML and so can you

NoSQL databases used to seem weird to some of us.

Distributed applications seemed odd at first.

Heck, Kubernetes....

(...nevermind, Kubernetes still seems pretty tough most of the time. But anyway...).

**ML systems are understandable and learnable.**



Photo by [C Dustin](#) on [Unsplash](#)

# The Future

7

# The Path Ahead:

## Predictions without Evidence



Photo by [Bruno Nascimento](#) on [Unsplash](#)

### Career changes

SREs will have to do ML  
(**ML Reliability Engineers?**)

but we will also need  
**data reliability engineers**

and model quality will have to be  
baked in to production  
engineering responsibilities.  
(Quality as the end-to-end SLO)

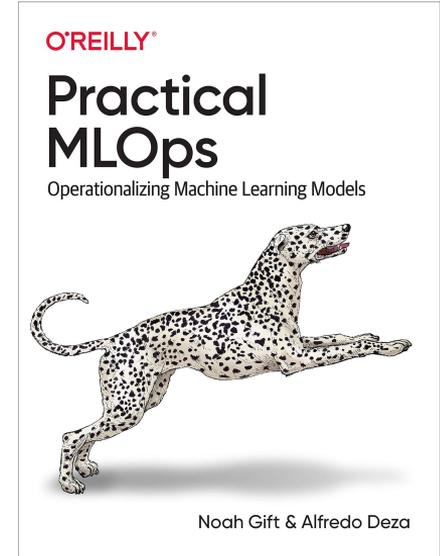
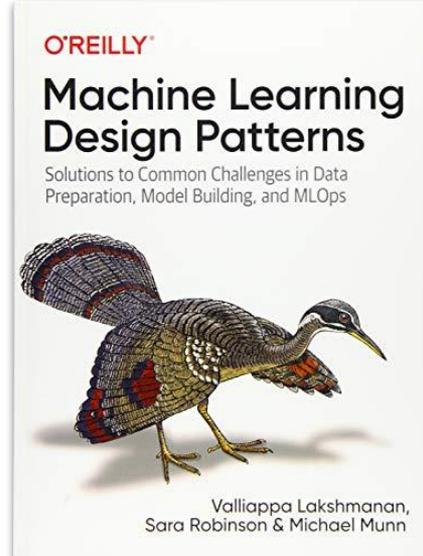
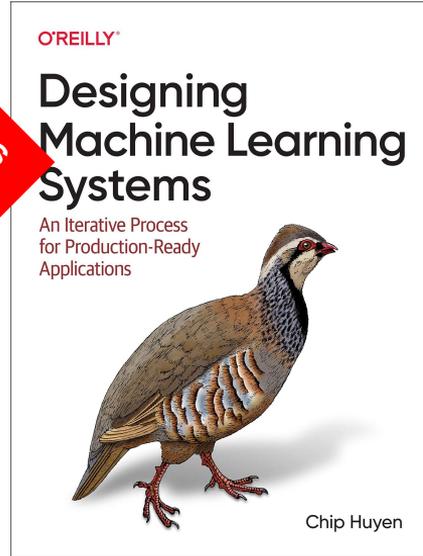
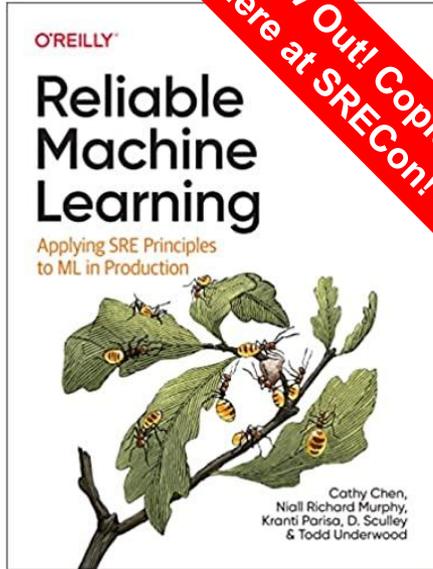
### Ethics

(Wildcard opinion!)

AI/ML raises huge and  
complicated ethical issues.  
SREs will be again (as we  
usually have been) at the  
forefront of addressing those  
issues as models enter into the  
world

# Possible Further Reading

Now Out! Copies  
here at SRECon!



Book covers copyright O'Reilly Media. Used with permission.

# Questions?



Photo by Pixabay from Pexels



**Thank you**